Draft version September 13, 2016
Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1609.03065v1 [astro-ph.IM] 10 Sep 2016

PROBABILISTIC CROSS-IDENTIFICATION IN CROWDED FIELDS AS AN ASSIGNMENT PROBLEM

Tamás Budavári[1,2,3] and Amitabh Basu[1]

Draft version September 13, 2016

ABSTRACT

One of the outstanding challenges of cross-identification is multiplicity: detections in crowded regions of the sky are often linked to more than one candidate associations of similar likelihoods. We map the resulting maximum likelihood partitioning to the fundamental assignment problem of discrete mathematics and efficiently solve the two-way catalog-level matching in the realm of combinatorial optimization using the so-called Hungarian algorithm. We introduce the method, demonstrate its performance in a mock universe where the true associations are known, and discuss the applicability of the new procedure to large surveys.

1. MOTIVATION

Modern astronomy surveys tenaciously observe the sky every night and their automated software pipelines produce a huge number of exposures. The datasets from separate telescopes and instruments are then routinely combined to enable all multicolor and time-domain studies. Over the last decade there has been a tremendous progress in the field of catalog matching to provide reliable resources for astronomical and cosmological measurements. Budavári and Szalay (2008) introduced a framework based on Bayesian hypothesis testing that properly incorporated astrometric uncertainties, which yielded superior results (Heinis *et al.* 2009), and proved to be flexible to accommodate a variety of scenarios from associating stars with unknown proper motions (Kerekes *et al.* 2010) or even models for radio morphology (Fan *et al.* 2013). These methods are also paired with efficient indexing ideas and search algorithms to make the process of matching fast on the largest catalogs, e.g., (Kunszt *et al.* 2001; Górski *et al.* 2005; Crittenden 2000; Gray *et al.* 2007; Lee and Budavári 2013)

The methods before, however, were only concerned with assigning a reliable quality measure of the association for a given set of sources. The likelihood of the association would be compared to the likelihood of the detections belonging to separate objects using the Bayes Factor and the previously used approaches do not resolve situations where a single detection appears in multiple associations with similar likelihoods. Previous recommendations to deal with such situations resorted to including more data, e.g., photometric measurements to refine the likelihood measurements (Marquez *et al.* 2014). In other words all candidate associations so far have been considered in isolation ignoring the fact that a star, for example, could only appear in a single match. Such exclusion rules can potentially affect a large fraction of the catalog especially in crowded fields.

In this paper we study catalog-level cross-identification to provide strategies for finding the most likely matched catalogs in which every association is valid and no detections appear in multiple matches. In Section 2 we formulate the catalog-level matching problem and map the likelihood optimization onto discrete minimization that is solved efficiently by the so-called Hungarian algorithm. Section 3 details our implementation and its application to realistic mock catalogs. In Section 4 we discuss the results and Section 5 concludes the paper.

Email: budavari@jhu.edu, basu.amitabh@jhu.edu
[1] Dept. of Applied Mathematics & Statistics, Johns Hopkins University, 3400 N. Charles St., MD 21218, USA
[2] Dept. of Computer Science, Johns Hopkins University, 3400 N. Charles St., MD 21218, USA
[3] Dept. of Physics & Astronomy, Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218, USA

2. CATALOG-LEVEL MATCHING

Matching entire catalogs is a much more challenging task than calculating the evidence for a match. Based on the same equations one has to find the optimal configuration such that all matches in the output catalog are globally optimal. Before we turn to the mathematical formulation, we introduce the notion of partitioning (Budavári and Loredo 2015) and illustrate the problem in a simple scenario.

Let us consider two catalogs with 2 sources in catalog 1 and only 1 in catalog 2. In this situation there are only 3 possibilities: the source in catalog 2 is either matched to one of the two sources in catalog 1 or it is a completely separate one by itself. If the pair (i, c) denotes item i in catalog c, we can write the possible partitions, excluding combinations of associations within each catalog, as

$$\mathcal{P}_0 = \Big\{\ \{(1,1)\},\ \{(2,1)\},\ \{(1,2)\}\ \Big\}$$
$$\mathcal{P}_1 = \Big\{\ \{(1,1)\},\ \{(2,1),(1,2)\}\ \Big\}$$
$$\mathcal{P}_2 = \Big\{\ \{(1,1),(1,2)\},\ \{(2,1)\}\ \Big\}$$

where $\mathcal{P}_0$ corresponds to three separate objects but $\mathcal{P}_1$ and $\mathcal{P}_2$ represent alternative matching scenarios with only two objects. It is true that, in general, different partitioning schemes can have a different number of associations or objects. We introduce some notation to make the problem mathematically precise. Let D denote the dataset of all (i, c) pairs for all catalogs c and items i in catalog c. A *valid partition* $\mathcal{P}$ is a set partition of D into disjoint subsets $S_o \subseteq S$ indexed by $o \in O_{\mathcal{P}}$ satisfying the condition that each set S_o contains at most one item from any catalog. More precisely, for two distinct tuples (i_1, c_1) and (i_2, c_2) in S_o, we must have $c_1 \neq c_2$. For example, with two catalogs, each S_o contains at most two tuples. An S_o of size 1 will be called an *orphan* and S_o

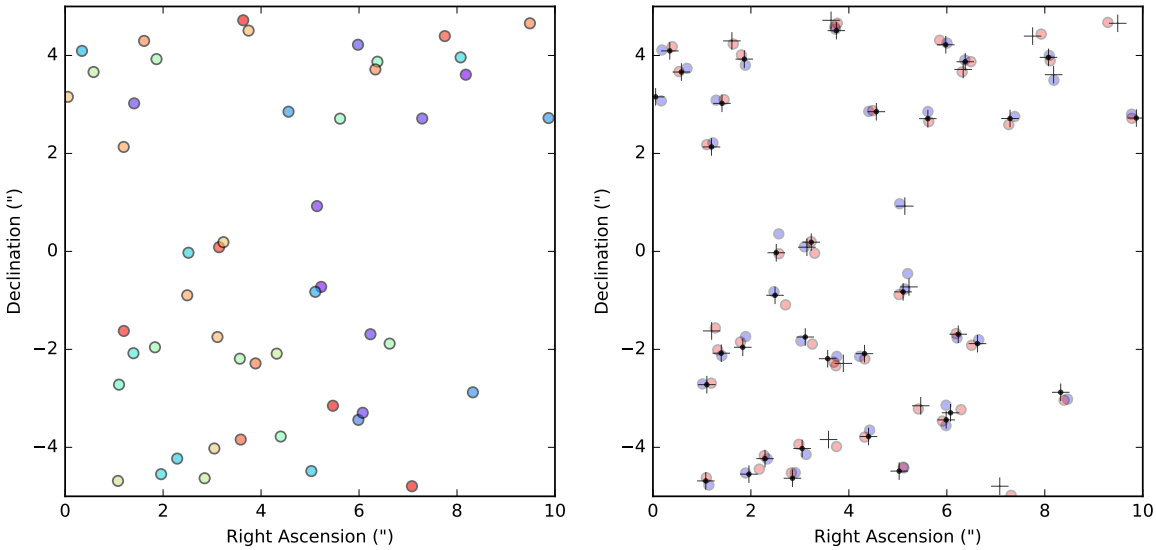

FIG. 1.— *Left:* A mock universe is constructed by randomly assigning coordinates (circles) and internal properties (color) to objects. *Right:* Simulated surveys observe different subsets of the objects with overlap. The + signs show the true directions of the objects but the *red* and *blue* dots illustrate the noisy measurements of the simulated sources. Objects that meet both selection functions are noted with small black dots on the true positions.

with size 2 will be called an *association*. The collection of indices $o \in O_{\mathcal{P}}$ will be called the set of *objects* of the valid partition $\mathcal{P}$.

Our goal is to find the best valid partition that maximizes the likelihood function of the valid partitions. The fundamental equations of the Bayesian cross-identification are still applicable to the individual associations but here we consider all combinatorial possibilities.

2.1. *Likelihoods of Matched Catalogs*

The likelihood of a given valid partition $\mathcal{P}$ is the probability density of the data $D=\{(i,c)\}$ given valid partition $\mathcal{P}$, and can be factored into the marginal likelihoods $\mathcal{M}_o$ of the objects $o \in O_{\mathcal{P}}$ of the valid partition,

$$\mathcal{L}(\mathcal{P}) \equiv p(D|\mathcal{P}) = \prod_{o \in O_{\mathcal{P}}} \mathcal{M}_o \tag{1}$$

The marginal likelihoods $\mathcal{M}_o$ are obtained by integrating the prior on the direction times the likelihood function over the entire sky (Budavári and Loredo 2015):

$$\mathcal{M}_o = \int d\boldsymbol{\omega} \;\; \rho_{\mathcal{C}(o)}(\boldsymbol{\omega}) \prod_{(i,c)\in S_o} \ell_{ic}(\boldsymbol{\omega}) \tag{2}$$

where $\ell_{ic}(\boldsymbol{\omega})$ is the likelihood function of source i in catalog c and $\rho_{\mathcal{C}(o)}(\boldsymbol{\omega})$ is the prior approriate for all catalogs $\mathcal{C}(o)$ that contribute to object o. For example, in the 2-way matching scenario there will only be up to two terms in the likelihood product that correspond to the two sources in the separate catalogs. The footprint of the catalogs along with the astrometric uncertainties determine the prior but in practice an isotropic prior is often an acceptable approximation for most associations when scaled (Budavári and Szalay 2008).

The marginal likelihood can be analytically calculated for the Fisher (1953) distribution, $\ell(\boldsymbol{\omega}) = f(\boldsymbol{x};\boldsymbol{\omega},\kappa)$, which is specified by two parameters: the concentration parameter κ, and a unit vector representing the direction of the mode, $\boldsymbol{\omega}$,

$$f(\boldsymbol{x};\boldsymbol{\omega},\kappa) = \frac{\kappa}{4\pi \, \sinh\kappa} \;\; \exp\Big(\kappa \, \boldsymbol{\omega}\boldsymbol{x}\Big) \tag{3}$$

In the limit of small uncertainties typical in today's astronomy surveys, this distribution on the surface of the sphere can be well-approximated with the bivariate normal distribution on the tangent plan.

Next we apply the usual logarithmic transformation and, instead of maximizing the likelihood, we minimize

$$A(\mathcal{P}) \equiv -\ln \mathcal{L}(\mathcal{P}) = -\sum_{o} \ln \mathcal{M}_o \tag{4}$$

When matching two catalogs, the terms in the sum correspond to pairs of sources in separate catalogs (defined as associations above), or orphan sources without a matching pair. We have a well-defined cost function to minimize by finding the best matching pairs. Similar problems have long been studied in discrete mathematics and algorithms are known to find the solution in polynomial time, that also run extremely efficiently in practice. We make a brief introduction of these ideas next, and show how to model the valid partition problem above using these ideas.

2.2. *The Assignment Problem*

One can pose the problem of minimizing (4) in the following equivalent way. Let the sizes of the two catalogs be n_1 and n_2. Introduce a $n_1 \times n_2$ matrix W, whose rows are indexed by the data points (i,c) in the first catalog and the columns are indexed by the data points in the second catalog.

The entry in the matrix W corresponding to a row indexed by (i_1,c_1) and a column indexed by (i_2,c_2) (where

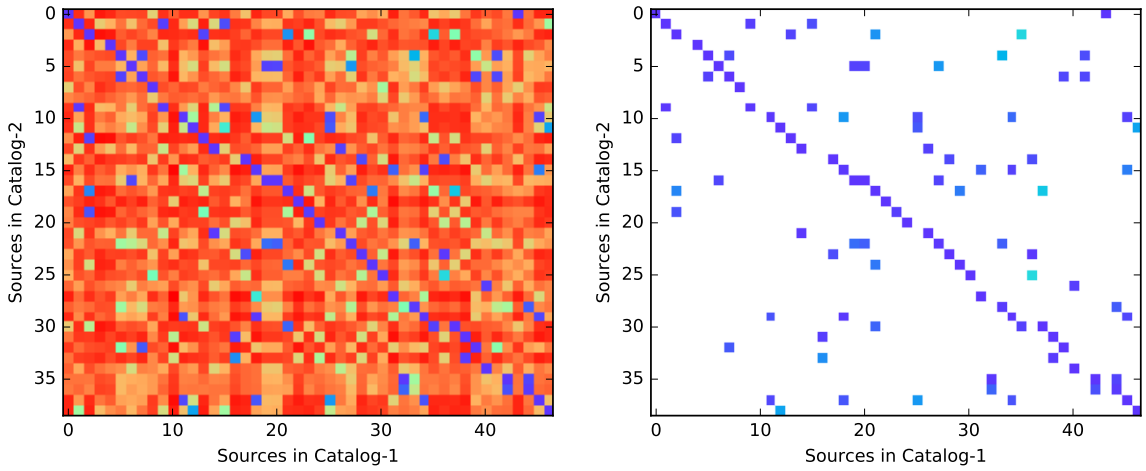

FIG. 2.— The weight matrix W that corresponds to the assignment problem contains the two catalogs along the axes. In the left panel, he *bluer* matrix elements indicate more likely possible associations and the *redder* ones are farther from each other. In the right panel, we plot the sparsified weight matrix where the very unlikely candidates are omitted; see text.

of course, by definition, c_1 is the label for the first catalog and c_2 is the label for the second catalog), is given by the formula:

$$-\ln\left(\int d\boldsymbol{\omega}\ \rho_{\{c_1,c_2\}}(\boldsymbol{\omega})\ \ell_{i_1c_1}(\boldsymbol{\omega})\ \ell_{i_2c_2}(\boldsymbol{\omega})\right)$$

Now here is the key observation. Let $\mathcal{U}$ be the set of matrices of size $n_1\times n_2$ with entries in $\{0,1\}$ such that every row and every column has at most one non zero entry. Any $P\in\mathcal{U}$ can be interpreted as a valid partition of D: For any 1 entry in the matrix, the objects corresponding to the row and column with this entry are matched and form an association. The remaining objects are all orphans; equivalently, if any object has only 0's in its corresponding row or column then it is assigned as an orphan. Conversely, for any valid partition $\mathcal{P}$ one can associate a matrix $P_{\mathcal{P}}\in\mathcal{U}$ that has the property that every non-zero entry corresponds to an association. Therefore, there is a one-to-one correspondence between valid partitions and $\mathcal{U}$. One can then show the following:

Theorem 2.1. *The valid partition $\mathcal{P}$ that minimizes* (4) *corresponds to a matrix $P\in\mathcal{U}$ that minimizes the trace* $\mathrm{Tr}\,(W^TP)$ *of the matrix W^TP. Thus, minimizing* (4) *is equivalent to solving*

$$\min\Big\{\mathrm{Tr}(W^TP) : P\in\mathcal{U}\Big\} \tag{5}$$

Given any matrix $W\in\mathbb{R}^{n\times n}$, the *Hungarian Algorithm* (Kuhn and Yaw 1955; Munkres 1957) is precisely a method to solve (5). This is a classically studied combinatorial optimization problem called the *Assignment problem*. The Hungarian Algorithm that solves it, has a guaranteed running time of $\mathcal{O}((n_1+n_2)^3)$, but performs extremely fast in practice.

3. APPLICATION

The new probabilistic catalog-level cross-identification approach was tested on synthetic catalogs of a mock universe. In such settings we can include realistic artifacts that match typical observations but still know the ground truth.

3.1. *Mock Objects and Surveys*

Our mock universe contains point-like objects in a small field of view. In addition to a random direction $\boldsymbol{x}$, each object has a random property $u_{01}\in[0,1]$, which could be thought of as the temperature of a star. The simulated sources are realizations of the objects with directions perturbed according to the chosen astrometric uncertainties. The selection effects, e.g., limiting magnitude, is determined by preset interval constraints on u_{01}. The surface density of a catalog can be adjusted by the length of the interval. The commonality of two catalogs is set by the overlap of the selection interval. Figure 1 illustrates a mock universe with the colors showing the properties of the objects on the left, and two catalog realizations are shown on the right.

The test scenario simulates catalogs in the flat-sky approximation using Gaussian with 0.1-arcsecond astrometric uncertainties, which is similar to that of the Sloan Digital Sky Survey (SDSS). Using an isotropic prior the weight matrix W is analytically calculated the same way as current procedures evaluate the Bayes factors. Figure 2 illustrates the structure and the values of the weights. The left and right panels differ only by the hopelessly bad matches (shown in red on the left), which can be eliminating upfront before even considering the combinatorial problem. In practice we always work with such sparse matrices of the candidates (right panel), which can be quickly calculated by existing procedures. We note the ridge of good matches in the original $n_1\times n_2$ matrix, which is an artifact of the simulations: sources appear

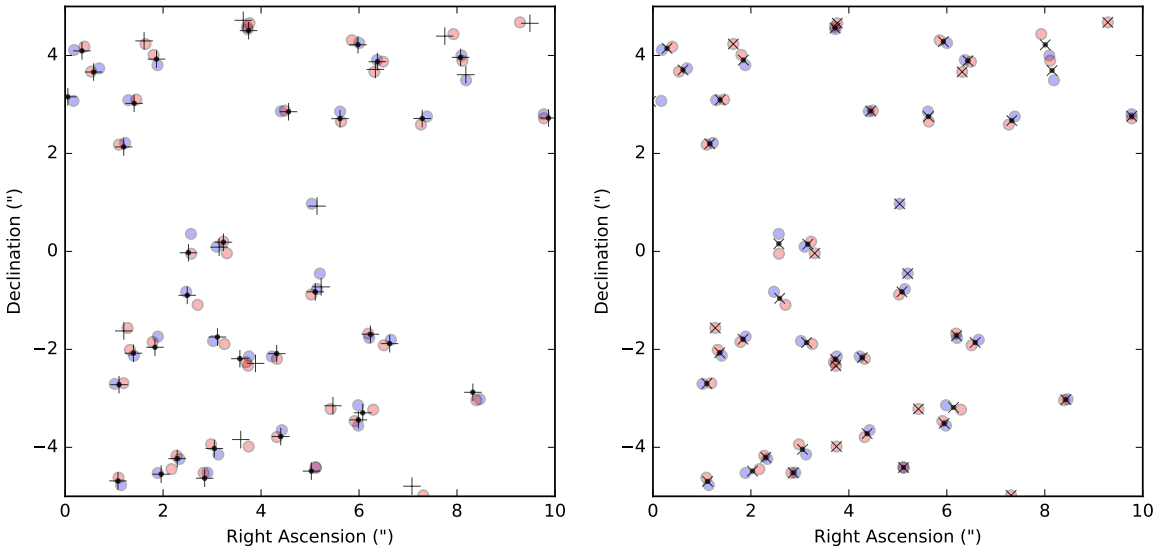

FIG. 3.— The two panels compare the truth in the *left* panel (same as in Figure 1) and the optimal assignment of the detections on the *right*, where the *crosses* show the best-guess direction of the associations. The only wrong association is in the top right corner, see text.

in the same order as the mock objects. The advantage of this view is that we see the large number of confused elements that do not follow the linear trend.

Figure 3 shows the ground truth (left) and the matched results (right). The + signs illustrate the directions of the true objects as before and the result of our matching procedure using the × glyphs (right). The simulated sources are displayed as red and blue dots in both panels for easy comparison. We see that even in this extremely confused field where the astrometric uncertainly would allow for multiple matches for almost all sources, the objects are perfectly resolved in almost all cases even in the most crowded parts of the field of view. The assignments correspond to objects with multiple detections and even the orphan sources are correctly isolated. The only exceptions are due to selection effects, as illustrated in the top right corner of the field. In the left panel we see that in the mock universe there are three separate objects: the middle one is detected in both catalogs but the other two are only in one or the other. In this situation with the coordinates so close to each other, the algorithm finds that instead of matching the middle sources to each other, it is better to assign them to the nearby sources. While this is incorrect, we see that given the data, the output shown in the right panel is the optimal solution: having 2 pairs is better than 1 pair and 2 orphans.

3.2. *Evaluation of Repeated Simulations*

Using the same simulation routines a suite of tests were performed to quantify the performance of the new approach in comparison to the traditional nearest neighbor methods, which in this simulation is equivalent to the maximum likelihood or Bayes factor matches due to the assumed homoscedastic astrometric errors.

For each mock we uniformly generate objects in a 3'×3' field of view with surface density of 400 per square arc minutes. Using an astrometric uncertainty of 0.04", two catalogs were generated including every object to best illustrate the effect of the exclusion rule.

The simulated catalogs were matched by both finding the closest sources and by solving the assignment problem with the Hungarian algorithm. Figure 4 summarizes the distribution of the number of incorrect matches in the two scenarios. The gray error bars represent the measurements and error estimated from the square root of the counts. We see that perfect matches practically never occur. The Poisson distribution with the parameter set to the sample mean (λ=7.9) is a close fit as shown by the solid bars.

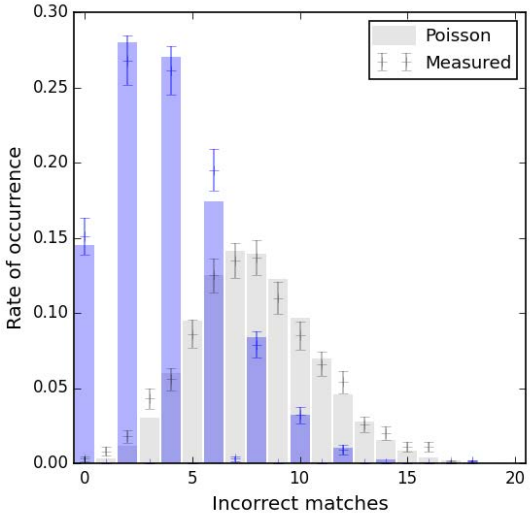

FIG. 4.— The rate of incorrect outputs when matching source lists repeatedly. The gray points shows the number of incorrect matches for the nearest neighbor associations in a 3'×3' field of view with surface density of 400 objects per sq. arcmin. The error bars simply represent the square root of the counts and the solid bars illustrate the Poisson distribution with λ set to the sample mean. The blue points and error bars correspond to results from our new approach, see text.

The blue error bars represent the new results. The first thing we notice is the gaps at odd numbers – mistakes tend to occur in pairs due to the exclusion rules. The new method provides perfect matches 15% of the time and the distribution is in general better skewed toward zero. While the traditional method yields more than 4 incorrect matches with 90% probability, the new method shrinks this to just 30%. The sample mean improves to 3.87 and the mode by a essentially a factor of 4. The shape of the distribution can be actually well approximated by a Poisson distribution with parameter set to the half the sample mean (λ=1.9) and stretched along the horizontal axis by a factor of two.

4. DISCUSSION

The results that correspond to the maximum likelihood solution of the entire catalog appear exceptionally reliable and go beyond the limitations of working with individual associations. With the marginal likelihoods for the winning associations we can go back and derive the probability of each match using the prior information on surface density of the sources in the catalogs Budavári and Szalay (2008).

4.1. *Computational Cost*

The proposed procedure is clearly more expensive than previous methods. As mentioned earlier, the general assignment problem solved by the Hungarian algorithm requires $\mathcal{O}((n_1+n_2)^3)$ time, where n_1 and n_2 are the number of rows and columns of the input matrix W (in our case $|D|=n_1+n_2$). For example, the 3600×3600 weight matrices of our simulations were processed by the solver called `linear_sum_assignment()` in the `optimize` module of SciPy (Jones *et al.* 2001) in approximately 2.2 seconds on a desktop computer. While this sounds prohibitively slow at first, when considering large astronomy surveys of hundreds of millions of sources, in practice the wall-clock time will be actually much shorter due to the sparsity of the weight matrix. In other words the list of candidate associations is actually only on the order of $\max(n_1, n_2)$ long instead of the number of pairs, which naively would be $\mathcal{O}(n_1 n_2)$.

In fact resolving the multiplicity issue can be reserved for the overlapping associations where friends-of-friends connected components emerge as problem cases. Luckily these can be dealt with separately, hence reducing the cardinality of the problem to much smaller numbers of sources. Budavári and Lubow (2012) used a greedy approach to find the best the partitions using clustering algorithms and Bayesian hypothesis testing. In the 2-way matching case, we can actually afford to do a global search for the optimum with existing methods such as the Hungarian algorithm.

4.2. *Likely Matched Catalogs*

Finding the optimal matched catalog is unfortunately not the final answer. There can be other realizations with the same or similarly good likelihoods. In principle these alternatives should be reported for further statistical analysis and marginalized over in the study.

While the approach applied in our study does not automatically provide the sub-optimal solutions, the optimum would a good starting point for some Markov chain Monte Carlo algorithm to explore alternative matched catalogs. The cost of storing or automatically generating these catalogs and providing the uncertainty estimates on the associations or their statistics needs to be studied further but is straightforward from the methodology's point of view.

5. SUMMARY

Our method builds on previous developments in the field of cross-identification. The candidate associations are found by efficient matching procedures using ideas in database indexing and search trees and their sparse weight matrix is derived from the Bayesian approach introduced earlier. The conceptual change is to do a joint optimization to find the best matched catalog among the possible competing scenarios.

We have demonstrated that the catalog-level matching problem can be efficiently solved in case of two input datasets. The maximum of the marginal likelihood of the entire matched catalog is found using discrete optimization such as the Hungarian algorithm in the case of our prototype. Despite the complexity of the problem, the sparse nature of rare $\mathcal{O}(n)$ associations suggests that the new method is applicable to the largest catalogs.

Matching more than two catalogs, however, remains to be a challenge as no analog of the Hungarian algorithm or similar exists for the general case. The only hope is that the associations become even sparser and in that limit new algorithms can be developed. Traditional matching procedures, which do not consider the multiplicity issue, can efficiently find candidate associations by adding catalogs one-by-one using a recursive formula (Budavári and Szalay 2008; Budavári *et al.* 2013). Our future plans include examining the connectivity graph of the associations and work with the problem cases separately.

ACKNOWLEDGEMENT

The authors are grateful to Tom Loredo and Steve Lubow for invaluable discussions on various aspects of this study. TB acknowledges partial support from the NSF via grant AST-1412566 and NASA via the awards NNG16PJ23C and STSci-49721 under NAS5-26555.

REFERENCES

T. Budavári and A. S. Szalay, ApJ **679**, 301 (2008), arXiv:0707.1611.

S. Heinis, T. Budavári, and A. S. Szalay, ApJ **705**, 739 (2009), arXiv:0910.3214.

G. Kerekes, T. Budavári, I. Csabai, A. J. Connolly, and A. S. Szalay, ApJ **719**, 59 (2010), arXiv:1006.2096 [astro-ph.GA].

D. Fan, T. Budavári, A. S. Szalay, C. Cui, and Y. Zhao, PASP **125**, 218 (2013), arXiv:1403.4358 [astro-ph.IM].

P. Z. Kunszt, A. S. Szalay, and A. R. Thakar, in *Mining the Sky*, edited by A. J. Banday, S. Zaroubi, and M. Bartelmann (2001) p. 631.

K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, ApJ **622**, 759 (2005), astro-ph/0409513.

R. G. Crittenden, Astrophysical Letters and Communications **37**, 377 (2000), astro-ph/9811273.

J. Gray, M. A. Nieto-Santisteban, and A. S. Szalay, eprint arXiv:cs/0701171 (2007), cs/0701171.
M. A. Lee and T. Budavári, in *Astronomical Data Analysis Software and Systems XXII*, Astronomical Society of the Pacific Conference Series, Vol. 475, edited by D. N. Friedel (2013) p. 235.
M. J. Marquez, T. Budavári, and L. M. Sarro, A&A **563**, A14 (2014).
T. Budavári and T. J. Loredo, Annual Review of Statistics and Its Application **2**, 113 (2015), http://dx.doi.org/10.1146/annurev-statistics-010814-020231.
R. Fisher, Proceedings of the Royal Society of London Series A **217**, 295 (1953).
H. W. Kuhn and B. Yaw, Naval Res. Logist. Quart , 83 (1955).
J. Munkres, "Algorithms for the assignment and transportation problems," (1957).
Z. Ivezic, T. Axelrod, W. N. Brandt, D. L. Burke, C. F. Claver, A. Connolly, K. H. Cook, P. Gee, D. K. Gilmore, S. H. Jacoby, R. L. Jones, S. M. Kahn, J. P. Kantor, V. V. Krabbendam, R. H. Lupton, D. G. Monet, P. A. Pinto, A. Saha, T. L. Schalk, D. P. Schneider, M. A. Strauss, C. W. Stubbs, D. Sweeney, A. Szalay, J. J. Thaler, J. A. Tyson, and LSST Collaboration, Serbian Astronomical Journal **176**, 1 (2008).
E. Jones, T. Oliphant, P. Peterson, *et al.*, "SciPy: Open source scientific tools for Python," (2001), [Online; accessed 2016-05-16].
T. Budavári and S. H. Lubow, ApJ **761**, 188 (2012), arXiv:1206.0644 [astro-ph.IM].
T. Budavári, L. Dobos, and A. S. Szalay, Computing in Science & Engineering **15**, 12 (2013)